DirectCertify

Unlock Your Potential, Certify Your Success

Cloudera

CDP-500 Exam

Cloudera Administrator Cloud Certification Exam

Exam Latest Version: 6.0

DEMO Version

Full Version Features:

• 90 Days Free Updates

• 30 Days Money Back Guarantee

• Instant Download Once Purchased

• 24 Hours Live Chat Support

Full version is available at link below with affordable price.

https://www.directcertify.com/cloudera/cdp-500

Question 1. (Multi Select)

Which of the following are critical prerequisites before onboarding a Cloudera Data Platform (CDP) environment on a public cloud provider (select all that apply)?

A: Properly configured VPC/VNet with appropriate subnets and routing.
B: A detailed disaster recovery plan.
C: Appropriate IAM roles/permissions for CDP to access cloud resources.
D: Implementation of a continuous integration/continuous delivery (CIICD) pipeline.
E: A well-defined security strategy, including network security, encryption, and access controls.

Correct Answer: A, C, E

Explanation:

Before deploying CDP, it's crucial to have the network infrastructure (VPC/VNet), cloud access permissions (IAM roles), and security strategy in place. While a DR plan and CI/CD pipeline are important, they are not strictly prerequisites for initial onboarding. They are important for operation but not specifically onboarding.

Question 2. (Single Select)

You are onboarding a new Cloudera Data Platform (CDP) environment in AWS. Your security team requires all network traffic to be encrypted both in transit and at rest. Which of the following options are BEST suited for achieving this goal with minimal administrative overhead?

A: Enable encryption at rest using AWS KMS managed keys for all S3 buckets and use TLS for all communication between CDP components. Configure all EC2 instances with AWS Systems Manager Agent for patching and vulnerability scanning.
B: Implement IPsec tunnels between all subnets in the VPC and use self-signed certificates for internal communication. Enable server-side encryption (SSE) with S3 managed keys (SSE-S3).
C: Utilize AWS CloudHSM for all key management and implement mutual TLS authentication for all CDP components. Enable S3 bucket encryption with client-side encryption using keys managed in CloudHSM.

D: Rely solely on the default encryption provided by S3 and EC2. Configure network ACLs for basic security.

E: Enable encryption at rest using AWS KMS Customer Managed Keys (CMK) for all S3 buckets and use TLS for all communication between CDP components. Additionally, implement security groups with the principle of least privilege for network access control.

Correct Answer: E

Explanation:

Using KMS CMKs gives you control and auditability over the encryption keys. TLS ensures encryption in transit. Security Groups based on the principle of least privilege provide robust network access control. Options B and C introduce unnecessary complexity. Option A uses AWS managed keys which give less control. Option D is insufficient for security.

Question 3. (Single Select)

You have a requirement to monitor the CPU utilization of all the worker nodes in your CDP cluster running on AWS. Which of the following monitoring tools, when integrated with Cloudera Manager, provides the MOST comprehensive and scalable solution?

A: CloudWatch
B: Ganglia
C: Nagios
D: Prometheus with Grafana
E: AWS x-Ray

Correct Answer: D

Explanation:

Prometheus, coupled with Grafana for visualization, offers a powerful and scalable solution for monitoring time-series data, including CPU utilization. It's well-suited for dynamic cloud environments and integrates well with containerized deployments (if applicable). CloudWatch is specific to AWS, making it not the best choice. Ganglia is older technology, and it doesn't scale as well. Nagios provides host and service monitoring but lacks the deep metrics analysis

capabilities. AWS X-Ray is a distributed tracing system, not for server level monitoring

Question 4. (Multi Select)

Which of the following are critical prerequisites before onboarding a Cloudera Data Platform (CDP) environment on a public cloud provider (select all that apply)?

A: Properly configured VPC/VNet with appropriate subnets and routing.
B: A detailed disaster recovery plan.
C: Appropriate IAM roles/permissions for CDP to access cloud resources.
D: Implementation of a continuous integration/continuous delivery (CIICD) pipeline.
E: A well-defined security strategy, including network security, encryption, and access controls.

Correct Answer: A, C, E

Explanation:

Before deploying CDP, it's crucial to have the network infrastructure (VPC/VNet), cloud access permissions (IAM roles), and security strategy in place. While a DR plan and CI/CD pipeline are important, they are not strictly prerequisites for initial onboarding. They are important for operation but not specifically onboarding.

Question 5. (Single Select)

You are configuring a Cloudera Data Engineering (CDE) service on Azure. You need to set the correct environment variables for Spark applications to access data stored in Azure Data Lake Storage Gen2 (ADLS Gen2). Which of the following environment variables (with their correct values) are REQUIRED for Spark to authenticate with ADLS Gen2 using a Service Principal?

A:

```
fs.azure.account.oauth2.client.endpoint=; fs.azure.account.oauth2.client.id=; fs.azure.account.oauth2.client.secret=
```

B:

```
dfs.adls.oauth2.access.token.provider.type=ClientCredential; dfs.adls.oauth2.client.id=; dfs.adls.oauth2.credential=; dfs.adls.oauth2.refresh.url=
```

C:

```
fs.azure.account.key..dfs.core.windows.net=
```

D:

```
spark.hadoop.fs.azure.account.oauth2.client.endpoint=; spark.hadoop.fs.azure.account.oauth2.client.id=;
spark.hadoop.fs.azure.account.oauth2.client.secret=; spark.hadoop.fs.azure.account.oauth2.msi.tenant=
```

E:

```
dfs.adls.oauth2.access.token.provider.type=ClientCredential; dfs.adls.oauth2.client.id=; dfs.adls.oauth2.credential=; dfs.adls.oauth2.refresh.url=;
spark.hadoop.fs.azure.account.oauth2.msi.tenant=
```

Correct Answer: D

```
spark.hadoop.fs.azure.account.oauth2.client.endpoint=; spark.hadoop.fs.azure.account.oauth2.client.id=;
spark.hadoop.fs.azure.account.oauth2.client.secret=; spark.hadoop.fs.azure.account.oauth2.msi.tenant=
```

Explanation:

Option D uses the correct 'spark.hadoop.fs.azure.account' prefix, which is necessary for Spark to recognize these configuration properties. It also includes the 'oauth2.msi.tenant' property for the tenant ID. The property is not required. Option C uses account key which is not optimal. Option B uses deprecated property names. Option A has incorrect property names.

Full version is available at link below with affordable price.

https://www.directcertify.com/cloudera/cdp-500

30% Discount Coupon Code: LimitedTime2025